JOB SHOP HEURISTICS AND STATISTICAL INFERENCE IN A TAXICAB SIMULATION

Tasker Generes
Armour and Company
Chicago, Illinois

A simulation of the taxi service system of the University of California Lawrence Radiation Laboratory (UCLRL) was undertaken to investigate (without disrupting normal operations) ways of modifying the operations of the existing system which would result in substantially better service at no great additional cost. The successful achievement of the above objective resulted from the recognition of the close relationship between job shop scheduling and the decision-making problem of UCLRL's taxi service system; and the applicability, therefore, of some of the rules and heuristics employed in job-shop scheduling. GPSS III was used to model the proposed system, rules and heuristics because it readily permitted the formulating of complex models. Although GPSS flexibility allowed accurate modeling of the majority of the system, GPSS approximations were still necessary to represent the distribution of arrivals. To defend the independent Poisson distribution a statistical inference study was used to establish that the model's distribution was robust with respect to the effects of the actual distribution of arrivals.

A brief description of UCLRL is necessary prior to any further discussion of how the simulation study was particularly applied.

UCLRL is located to the north of the Berkeley campus on a long strip of approximately 540 acres. Spread out over this large area are more than fifty buildings with over 3,000 employees (see map). The scattering of many buildings over a wide area together with the many steep hills and winding roads make walking time consuming and very inconvenient. Due to the nature of their work, scientists and engineers are often going from one building to another.

Historically the transportation problem was attacked by the use of four buses on prescribed routes. This arrangement soon became inadequate and a system of four radio-dispatched taxis was installed. However, there was still a frustrating delay in service which if reduced would be extremely valuable to UCLRL. Before any further investment was made, management was convinced that a simulation system study should be carried out to determine the appropriate responses, if any, that could be made to improve the transportation situation.

The possible alternatives appeared to be: to increase the number of cabs, to increase the capacity of the cabs, and to change the dispatch rules. The first two alternatives obviously would have involved added expense, which was considered undesirable. Further, it seemed that little could be done to improve upon the monitor's present performance. It should be noted here that the monitor's decision-making was seen to be analogous to that in job-shop scheduling if we were to think of taxis as machines, a call for a taxi as a job, and the request for transportation between two specific locations as characterizing the nature of the job. This isomorphism suggested that possibly some of the rules and heuristics which had been developed in job shop scheduling might be applied. Although one would immediately think of the various studies of priority rules at a detailed scheduling level, it didn't seem that the solution of the problem was in that direction, but on a more general level. It was, and is, generally recognized that it is desirable to have uniform jobs, since imposing such uniformity will reduce the amount of extra effort required of the shop in coping with the unique characteristics of specific jobs.

In the context of the taxi problem, reducing the number of possible jobs is equivalent to reducing the set of possible embarkation and debarkation points. This can be accomplished by designating a group of buildings as a zone and requiring all departures from and arrivals to a zone at a particular spot. Thus zones would represent a partitioning of the many buildings in the existing system into non-overlapping areas such that all those leaving any building in a zone will be picked up at a single depot within the zone, and all those arriving at any building will be dropped off at the depot.

It was conjectured, therefore, that zoning would improve the taxi service. The objective of the simulation became then to test the effects of zoning and of varying the capacity and number of cabs. Fourteen zones were chosen. The bases for choosing were: 1) the buildings within each zone must be within close proximity and easy walking distance of one another, 2) each zone should have a sufficiently high rate of calls for taxis, and 3) every building must be in only one zone. All the necessary information was available for these decisions to be made since the monitor kept a daily log of the time, origination and destination points of each call for service.

In formulating the model, it was assumed that the inter-departure times of individuals between all pairs of zones were independent and exponentially distributed with mean equal to the observed mean interdeparture times. The actual distributions could not accurately be described as such; however, I hoped that the assumptions would be a fair approximation to the actual distributions. A simulation run could then have been made to test the model in which trips were generated at exactly the same times as they were during the period which was simulated. However, this would have involved the use of a function beyond the capacity of GPSS. The quandary was obvious.

The solution was a statistical inference study to determine if the Poisson distribution was a robust description of the underlying distribution for queue arrivals.

Let D_1 = the underlying distribution, D_2 = our distribution, R_1 = the rules for taxi service as they now exist, R_2 = the rules with zoning, and $F(R,D)$ = the system service when the trips are from D and rules from R are used. The hypothesis that $F(R_2,D_2)$ is a good estimate of $F(R_2,D_1)$ is tested by using the methodology of statistical inference and the tool of simulation. The natural approach is to take a "sample" to obtain information which would enable an inference to be made with some "reasonable" degree of certainty as to whether or not our hypothesis is correct. A simulation was done, therefore, generating samples of $F(R_2,D_2)$ and $F(R_2,D_1)$ for relatively short time periods instead of the full length of the main simulation. The inference study seemed to support hypothesis although the sample was necessarily small.

Having established some confidence in the results of the model both from the above approach and from the reasonableness of the results themselves, the next question is how does the model function? The model is simple in concept but fairly intricate in actuality.

The operation of the program is, basically, as follows: Passenger arrivals are placed in chains, passengers are released from these queues and placed in cab chains according to decisions made by the control transactions, and the passengers are kept in the cab until they are released by the control.

Pertinent statistics are maintained throughout. In combination with the economics of the system these statistics (see appendix) show that the optimum system for UCLRL is zoning with five taxi cabs in service. Apart from this practical result for UCLRL, I believe the study illustrates a good problem solving approach which is useful in situations that seem intractable, i.e. cast the problem into a different context. Hopefully the transformed problem may be one which has a recognized solution in the new space. If this desirable result isn't immediate, it may yet be beneficial because of familiarity with the body of knowledge and tools available for use in the new space. Another useful method illustrated is the use of the simulation in two ways - first, for an inference study to establish confidence in the model's results and, second, in the simulation of the model itself. Sampling methods such as these seem to be a technique easily applied to many situations to provide another measure of the model validity.